

AD-A034 598

DECISIONS AND DESIGNS INC MCLEAN VA
INTERACTIVE GRAPHICAL AIDS FOR BAYESIAN HIERARCHICAL INFERENCE. (U)

F30602-76-C-0156

F/G 12/1

DEC 76 S BARCLAY

UNCLASSIFIED

TR-76-9

RADC-TR-76-308


NL

| OF |
AD
A034598



END

DATE
FILMED
2-77



ADA034598

J
12

RADC-TR-76-308
Final Technical Report
December 1976



INTERACTIVE GRAPHIC AIDS FOR BAYESIAN HIERARCHICAL INFERENCE

Decisions and Designs, Incorporated

Approved for public release;
distribution unlimited.


J
D D C
REF ID: A65116
JAN 21 1977
D

ROME AIR DEVELOPMENT CENTER
AIR FORCE SYSTEMS COMMAND
GRIFFISS AIR FORCE BASE, NEW YORK 13441

This report has been reviewed by the RADC Information Office (OI) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public including foreign nations.


This report has been reviewed and is approved for publication.

APPROVED:




JOHN J. ATKINSON
Project Engineer

APPROVED:



HOWARD DAVIS
Technical Director
Intelligence and Reconnaissance Division

FOR THE COMMANDER:



JOHN P. HUSS
Acting Chief, Plans Office

62702F

Do not return this copy. Retain or destroy.

MISSION
of
Rome Air Development Center

RADC plans and conducts research, exploratory and advanced development programs in command, control, and communications (C^3) activities, and in the C^3 areas of information sciences and intelligence. The principal technical mission areas are communications, electromagnetic guidance and control, surveillance of ground and aerospace objects, intelligence data collection and handling, information system technology, ionospheric propagation, solid state sciences, microwave physics and electronic reliability, maintainability and compatibility.



UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)


(19) REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
(18) 1. REPORT NUMBER RADC-TR-76-308	2. GOVT ACCESSION NO. (9)	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) INTERACTIVE GRAPHICAL AIDS FOR BAYESIAN HIERARCHICAL INFERENCE	5. TYPE OF REPORT & PERIOD COVERED Final Technical Report	
(10) 7. AUTHOR(s) Scott Barclay	6. PERFORMING ORG. REPORT NUMBER TR-76-9	
8. PERFORMING ORGANIZATION NAME AND ADDRESS Decisions and Designs, Incorporated 8400 Westpark Drive/Suite 600 McLean VA 22101	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 4594 1222	
11. CONTROLLING OFFICE NAME AND ADDRESS Rome Air Development Center (IRDA) Griffiss AFB NY 13441	(11) 12. REPORT DATE December 1976	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Same (12) 86P.	13. NUMBER OF PAGES 78	
	15. SECURITY CLASS. (of this report) UNCLASSIFIED	
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE N/A	
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.	(16) 4594 (17) 12	
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) Same		
18. SUPPLEMENTARY NOTES RADC Project Engineer: John J. Atkinson (IRDA)		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Bayesian Hierarchical Inference Likelihood Ratio Likelihood Distribution Intervening Variable Bayes' Theorem		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) Bayesian hierarchical inference is an optimum methodology for diagnosing the probable state of an unknown variable from data which are known, but which are known, but which are only indirectly related to the variable of interest. The present report describes a general procedure for Bayesian hierarchical inference implemented on an interactive graphics computer. It allows an intelligence analyst to build a hierarchical inference tree structure, assess the required probability distributions, and compute the likelihood distribution over the	next page JB	

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

target variable. The procedure allows variables to be either, discrete, such as the type of research being pursued by a foreign government, or continuous, such as the personnel strength of a foreign tank regiment.

ACCESSION FOR	
NTIS	White Section <input checked="" type="checkbox"/>
DDO	Blue Section <input type="checkbox"/>
UNANNOUNCED	<input type="checkbox"/>
NOTIFICATION	
BY	
DISTRIBUTION/AVAILABILITY CODES	
ESTL.	AVAIL. AND/OR SPECIAL
A	



UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

SUMMARY

Interactive Graphical Aids For Bayesian Hierarchical Inference

1.0 The Technical Problem

The solution of many intelligence problems requires the use of data that are only indirectly related to the events of primary interest to the intelligence analyst. The analyst must make use of chained or cascaded inferences of the form: "the observation of event A makes event B more likely and the increased likelihood of B makes C (the event of concern) more likely." This form of inferencing, which arises in the solution of such diverse problems as indications and warning, and order-of-battle, is very difficult for the unaided analyst to handle intuitively. A formal solution to this problem, a general Bayesian hierarchical inference model, was developed under an earlier RADC contract. Although pilot applications of the methodology in the intelligence community were promising, it was obvious that widespread acceptance of the approach would depend upon the development of highly user-oriented computer implementations that could provide analysts a natural access to the methodology without requiring them to become mathematicians in the process. The object of the current contract was to develop a design concept for such an analyst aid and to develop and implement key software modules identified in the design concept.

2.0 Technical Background

Bayesian hierarchical inference is an optimum methodology for diagnosing the probable state of an unknown variable from data which are known, but which are only indirectly related to the variable of interest. The basic scheme is to build a structure from the top down, that is, from the target variable at the top, through a hierarchy of intervening variables, to the known data at the bottom. The links between variables in this structure are defined in terms of likelihood distributions, the relative likelihood of each of the states of the lower-level variable given each state of the upper level variable.

The methodology aggregates all of these likelihood distributions into a composite likelihood distribution; one which incorporates all of the observed data. This distribution is then combined with the probability distribution representing the analyst's prior opinion using Bayes' Theorem to obtain the probability distribution over the target variable given all the data.

3.0 Accomplishments

In order to make Bayesian hierarchical inferences practical, a procedure was designed which not only carries out the mathematical computations implied by the methodology, but also allows the user to structure his model quickly and easily, and to assess the requisite likelihood distributions with a minimum of effort, in a fashion which is natural to him. The procedure is applicable to a wide range of intelligence problems, incorporating both discrete and continuous variables at all levels in the model structure.

A key part in the implementation of this design was the development of five user interface software modules. These modules not only display the tree structure and probability distributions by drawing them on a display screen, but they also allow the user to interact in real time by using a light pen to select new options, move probability density functions, or adjust the heights of probability histograms. This technical report describes the overall design concept and illustrates how the user interface modules would be used by a hypothetical intelligence analyst solving an order-of-battle problem.

4.0 Implications For Further Research

The next phase in the research should be to implement the complete design concept by linking together the user interface modules, and adding the computational rollback algorithms. Once the design concept has been fully implemented and tested, the complete procedure should be tested in an operational environment using DIA analysts. Two promising areas of application are indications and warning, and order-of-battle.

EVALUATION

The results of this technical effort will be extremely valuable to the Intelligence Information Handling Projects for which RADC is responsible. The methodology and the techniques developed have clearly shown that substantive intelligence analysts can employ Bayesian Hierarchical Inference with these user-oriented graphical aids. Moreover, the procedure is applicable to a wide range of intelligence evaluation/assessment problems.

The results from this effort conducted under TPO 4, Intelligence Data Handling, will be incorporated into several of RADC Engineering Development programs where decision analysis aids are required. Initial introduction of the decision analytic techniques has occurred in HQ EUCOM and HQ PACOM where they have been applied to Warsaw PACT Order of Battle and Korean Indications and Warning respectively.



JOHN J. ATKINSON
Project Engineer

TABLE OF CONTENTS

	<u>Page</u>
1.0 INTRODUCTION	1
2.0 STRUCTURING THE MODEL	3
3.0 ASSESSMENT OF LIKELIHOOD DISTRIBUTIONS	17
3.1 CASE 1: Discrete Upper and Lower Variables	17
3.2 CASE 2: Discrete Upper Variable, Continuous Lower Variable	28
3.3 CASE 3: Continuous Upper Variable, Discrete Lower Variable	33
3.4 CASE 4: Continuous Upper and Lower Variables	55
4.0 PROCESSING THE MODEL	72

1.0 INTRODUCTION

Bayesian hierarchical inference is a methodology for revising prior opinion in the light of information which is only indirectly related to the variable of interest. It is essentially a diagnostic tool; that is, it is used to diagnose the current but unknown state of a variable, rather than to predict the unknown future state of a variable. There are two basic processes involved in Bayesian hierarchical inference: the first involves the computation of a composite likelihood ratio based on a number of pieces of information, and the second involves the revision of the prior probability distribution over the target variable using the composite likelihood ratio. The latter process is simply Bayes' Theorem; it is independent of the process by which the likelihood ratio was generated.

The former process, however, is the essence of all hierarchical inference. It arises because the data which the inference maker has at hand are only indirectly related to the variable of interest. The basic scheme is to build a structure from the target variable at the top to the data at the bottom, using a hierarchy of intervening variables. While the true states of these intervening variables, by definition, cannot be known for certain, the data provide information which makes some of the states of the lowest level intervening variables more or less likely; and the probable states of these intervening variables impact the probability of other intervening variables, and so on up through the hierarchy until the uppermost level, the variable of interest, is reached.

In practice, the usual difficulty in applying the methodology is that the mathematics of combining the implied probability distributions of all of the data into a composite likelihood distribution is quite tedious and difficult without a computer. In addition, structuring even a small model requires the assessment of quite a large number of probabilities, a process which can be considerably enhanced using computer-based assessment techniques. In particular, the use of interactive graphics can provide rapid and efficient assessments, with a variety of methods of feedback.

For some problems, moreover, the variables of interest are defined on a continuous scale. For example, one might be interested in the probable speed of a new airplane, the range of a submarine, or the number of personnel stationed in a particular garrison. In these cases it is almost hopeless to try to elicit the required likelihood distributions, display the implications of a user's assessments, and compute the

desired probability distribution over the variable of interest without the aid of an interactive graphics computer.

In the description which follows, the reader is led through a hypothetical session with an intelligence analyst who is assessing the probable number of personnel associated with a particular foreign division. In this session the analyst will structure his model, assess four different types of likelihood distributions, and obtain the composite likelihood distribution for the entire set of data. In addition, the analyst will inspect the contribution of several items of data and intervening variables to the aggregate likelihood distribution.

2.0 STRUCTURING THE MODEL

At the beginning of the procedure, the analyst is shown a menu which requests that he specify whether he wishes to create a new model or process an old model [Figure 2-1]. In this instance he wants to create a model, so he touches the words NEW MODEL with the light pen and is given a new menu of options [Figure 2-2]. This basic menu allows the analyst to 1) specify the structure of the model by creating, deleting, connecting, or disconnecting nodes to form a hierarchical tree structure; 2) assess the likelihood distributions which represent the connections between nodes in the tree; 3) examine previously assessed likelihood distributions or distributions created during the processing of the tree; 4) save the model for future use; and 5) process the tree to obtain the desired likelihood distribution.

To begin structuring the tree, the analyst touches the light pen to CREATE NODE, and the program responds by requesting the name he wishes to assign to the new node [Figure 2-3, line 1]. The analyst wishes to create a node to represent the variable "personnel strength," so he types PERS [line 2]. The program next requests the number of states of the variable PERS [line 3]. Since personnel strength is a continuous variable, the analyst types the letter "C" [line 4]. At this point, the program displays the node in the center of the screen [Figure 2-4]. The program displays all created nodes in the center of the screen, but they are movable so the analyst may put them wherever he wishes. Since he wants to locate PERS at the top of the tree, the analyst moves it upward by "pulling" it with the light pen [Figure 2-5]. In a similar fashion, the analyst creates the next node, representing readiness category. He selects CREATE NODE, the program requests the name [Figure 2-3, line 5], he types CAT [line 6], the program requests the number of states [line 7], and the analyst states that there are three readiness categories [line 8]. Having received the required information, the program displays the CAT node in the middle of the screen [Figure 2-6], and the analyst moves it up [Figure 2-7]. At this point, the analyst wants to establish the connection between CAT and PERS, so he touches the CONNECT menu option. The program waits until he designates which nodes are to be connected, and in which order. In this case, the information about readiness category will help to determine the personnel strength, so the analyst touches first the CAT node and then the PERS node to indicate that CAT impacts PERS. The program connects them in the designated direction with an arrow [Figure 2-8].

The second source of information about personnel strength comes from the capacity of the division's barracks. The




Figure 2-1




Figure 2-2

1 LABEL FOR NODE?
2 PERS

3 NUMBER OF STATES?
4 0

5 LABEL FOR NODE?
6 CAT

7 NUMBER OF STATES?
8 3

9 LABEL FOR NODE?
10 BARCAP

11 NUMBER OF STATES?
12 0

13 LABEL FOR NODE?
14 EQUIP

15 NUMBER OF STATES?
16 0

17 LABEL FOR NODE?
18 AREA

19 NUMBER OF STATES?
20 0

21 LABEL FOR NODE?
22 COUNT

23 NUMBER OF STATES?
24 1

25 LABEL FOR NODE?
26 HUMINT

27 NUMBER OF STATES?
28 1

29 LABEL FOR NODE?
30 DIST

31 NUMBER OF STATES?
32 2

33 LABEL FOR NODE?
34 ACTIVE

Figure 2-3

35 NUMBER OF STATES?
36 1

37 LABEL FOR NODE?
38 UNITS

39 NUMBER OF STATES?
40 1

41 LABEL FOR NODE?
42 OBS

43 NUMBER OF STATES?
44 1

Figure 2-3 (Continued)




Figure 2-4




Figure 2-5




Figure 2-6




Figure 2-7




Figure 2-8

analyst therefore creates a node representing barracks capacity, BARCAP [Figure 2-3, lines 9 and 10], which is a continuous variable [lines 11 and 12]. The program displays the node [Figure 2-9], the analyst moves it [Figure 2-10], and connects it to PERS [Figure 2-11].

This completes the specification of the second level of the tree. If the analyst knew for certain which category the division was in, and knew the barracks capacity, he could discover the likelihood distribution over personnel strength by assessing the likelihood distribution for category given personnel strength, and the likelihood distribution for barracks capacity given personnel strength, and processing the model. However, in this case the analyst is not certain about the division's readiness category. One of the indicators of category, however, is the percentage of full-strength equipment the division has. Accordingly, the analyst defines a new variable, EQUIP [Figure 2-3, lines 13 and 14], which is a continuous variable [lines 15 and 16]. The program displays the node [Figure 2-12], the analyst moves it under CAT [Figure 2-13], and connects it [Figure 2-14]. Similarly, since the analyst does not know the barracks capacity of the division but does know something about the area of the barracks, he creates a continuous node called AREA [Figure 2-3, lines 17 through 20], which the program then displays [Figure 2-15]. He moves it under BARCAP [Figure 2-16], and connects it to BARCAP [Figure 2-17].

The analyst continues in this fashion, adding nodes under the previously created nodes until he reaches nodes whose state he knows for certain or will know for certain, called data nodes. Data nodes are handled a little differently from other nodes by the program. When asked to process the tree structure, the program will recognize data nodes because they do not have any arrows pointing to them. If a data node has more than one state, the user will be asked which state has occurred. This allows the analyst to build a general model for a recurring situation and to specify the current state of data whenever the model is processed.

The data nodes include COUNT, the observed equipment count for the division, HUMINT, the report of a soldier previously assigned to the division, DIST, the distance of the division from the border, ACTIVE, the activity level of the division, and UNITS, the observed subordinate units attached to the division [Figure 2-3, lines 21 through 40, Figure 2-18]. At this point, the analyst is just about to add a data node which provides information about barracks area, but the displayed structure is a bit crowded. He can get around this difficulty simply by moving nodes around with the light pen until they are situated more conveniently. The analyst moves all of the nodes




Figure 2-9




Figure 2-10




Figure 2-11




Figure 2-12




Figure 2-13




Figure 2-14




Figure 2-15




Figure 2-16




Figure 2-17




Figure 2-18

[Figure 2-19], then creates the new node, OBS [Figure 2-3, lines 41 and 42], which stands for the observed barracks area. Since he knows the area he has observed, he specifies that OBS has one state [lines 43 and 44]; the new node is created, and the analyst moves it under AREA and connects it [Figure 2-20].




Figure 2-19




Figure 2-20

3.0 ASSESSMENT OF LIKELIHOOD DISTRIBUTIONS

At this point the analyst has specified the complete structure relating all the data he has to the target variable, personnel strength. Before processing the structure, however, he must assess the likelihood distributions which form the connections among the variables in the model. He can do this by touching the light pen to the ASSESS option, then touching the particular connection which represents the assessment he wishes to make. However, the response the program then takes will depend upon the states of the two variables connected.

The procedure distinguishes four different types of assessments: 1) the lower node and upper node are both discrete; 2) the upper node is discrete, but the lower node is continuous; 3) the upper node is continuous while the lower node is discrete; and 4) both nodes are continuous. The section which follows describes each of these four types of assessments as the analyst specified the likelihood distributions for the model he has created.

3.1 CASE 1: Discrete Upper and Lower Variables

The simplest case of a likelihood distribution occurs when both the upper variable and the lower level variable are discrete. In this example, the analyst has selected the connection between DIST, the distance of the division from the border, and CAT, the readiness category as his first assessment. The program first prints out the names of the two nodes, and the number of states associated with each [Figure 3-1, lines 1 through 4]. Next the program displays a probability histogram [Figure 3-2]. Since the complete likelihood distribution will be composed of a separate probability distribution for each level of the upper variable, the program elicits the assessments for each of the levels of the upper variable, starting with level 1. Since the variable DIST has two states, there are two bars on the probability histogram, which the program has set initially to the same height.

The analyst believes that if the division is in category 1, that is, if variable CAT is at level 1, then it is more likely that the division is near the border, rather than far from the border, that is, that the variable DIST is at level 1 rather than at level 2. In fact, the analyst believes, as a first approximation, that it is about twice as likely to be near rather than far. So he adjusts the probability histogram by touching the top of the bar whose height he wishes to change and by moving it up or down. In this example, he moves the top of the right bar down until it is about half as high as the bar on the left, to about 0.25 [Figure 3-3]. However, since the two states of DIST represent a mutually exclusive and

1 THE UPPER NODE IS "CAT".
2 IT HAS 3 STATES.

3 THE LOWER NODE IS "DIST".
4 IT HAS 2 STATES.

5 NEW VALUE?
6 0.15

7 NEW VALUE?
8 0.85

9 NEW VALUE?
10 0.16667

Figure 3-1




Figure 3-2




Figure 3-3

exhaustive set, their probabilities must sum to 1.0. The analyst touches NORMALIZE with the light pen, and the program normalizes the distribution by keeping the ratio between the two probabilities the same, but making sure the probabilities sum to 1.0 [Figure 3-4]. Satisfied that this is a good initial representation, he touches RETURN with the light pen, and the program proceeds to the next assessment [Figure 3-5]. At this point, the analyst is unsure of the distribution when the division is in category 2, so he does not touch the histogram, but proceeds by touching RETURN. The program presents the display for level 3 [Figure 3-6]. Here the analyst believes that it is very unlikely that a category 3 division would be near the border, so he decreases the probability of a distance near the border, and increases the probability of a longer distance [Figure 3-7]. Finished with this assessment for the moment, the analyst touches RETURN.

The program now displays the complete distribution in matrix form [Figure 3-8]. The analyst decides that a figure of .15 is about right for the probability that a category 3 division will be near the border, but wishes to change the probabilities to round numbers. To change any probability in the matrix, the analyst must touch it with the light pen. He touches the number .145 with the light pen, it disappears from the screen [Figure 3-9], and the program requests the new value [Figure 3-1, line 5]. He enters .15 [line 6], and the program displays the new probability [Figure 3-10]. The analyst now wishes to change the number .855 to .85. He does this because he knows that the columns in the matrix must be normalized before they are used in computing the composite likelihood distribution. If he does not change the .855 probability, the program will normalize the column, and neither of the probabilities will be round numbers. He touches the number, it disappears from the screen [Figure 3-11], the program requests the new number [Figure 3-1, line 7], he types it in [line 8], and the new number is displayed [Figure 3-12].

Having made this assessment, the analyst now believes that if the division were in category 1, it would be a little more likely to be near the border than he originally thought. He wishes to use the probability histogram to reassess the probabilities, so he touches the column heading for the first column. The program displays the probability histogram reflecting his current assignment of probabilities [Figure 3-13]. The analyst now believes that it is three times as likely for a category 1 division to be near the border rather than far from the border, so he increases the probability of state 1 to .9, and then decreases the probability of state 2 to .3 [Figure 3-14] to represent the ratio 3:1. He touches NORMALIZE, the program normalizes the probabilities [Figure 3-15], he touches RETURN, and the program returns to the matrix [Figure 3-16].

UPPER NODE: CAT
LOWER NODE: DIST
CURRENT STATE OF UPPER NODE: 1

NORMALIZE
RETURN




Figure 3-4

UPPER NODE: CAT
LOWER NODE: DIST
CURRENT STATE OF UPPER NODE: 2

NORMALIZE
RETURN




Figure 3-5




Figure 3-6




Figure 3-7

NORMALIZE
FILE DISTRIBUTION

UPPER NODE
CAT

	1	2	3
LOWER NODE	1	.663 .500 .145	
DIST	2	.337 .500 .855	

Figure 3-8

NORMALIZE
FILE DISTRIBUTION

UPPER NODE
CAT

	1	2	3
LOWER NODE	1	.663 .500	
DIST	2	.337 .500 .855	

Figure 3-9




Figure 3-10




Figure 3-11




Figure 3-12




Figure 3-13




Figure 3-14




Figure 3-15

NORMALIZE
FILE DISTRIBUTION

UPPER NODE
CAT

	1	2	3	
LOOKER	1	.750	.500	.150
DIST	2	.250	.500	.850

Figure 3-16

NORMALIZE
FILE DISTRIBUTION

UPPER NODE
CAT

	1	2	3	
LOOKER	1	.750	.500	.150
DIST	2	.250		.850

Figure 3-17

Next, the analyst decides that the probability distribution associated with a category 2 division is the same as the probability distribution associated with a category 1 division. In other words, the two types of divisions are equally likely to be located near the border. In this case, rather than changing both numbers in the matrix, the analyst decides to change one of them so that their ratio is 3:1. He touches the lower probability in column 2 with the light pen, it disappears [Figure 3-17]; the program requests a new value [Figure 3-1, line 9], he types in .16667, which is a third of .500 [line 10], and the new number is displayed [Figure 3-18]. Now the analyst touches NORMALIZE with the light pen, and all columns in the matrix are normalized [Figure 3-19]. Having completed his assessment of this distribution, the analyst touches FILE DISTRIBUTION, and the program stores his assessments and returns to the tree structure diagram [Figure 3-20].

3.2 CASE 2: Discrete Upper Variable, Continuous Lower Variable

The analyst next touches ASSESS and selects the connection between CAT and EQUIP. The program, as usual, begins the assessment of the likelihood distribution by verifying the identity of the variables and the number of states associated with each [Figure 3-21, lines 1 through 4]. The program proceeds by obtaining a likelihood distribution for each state of the upper-level variable. To define each distribution, the program requires the minimum value of EQUIP when CAT is in state 1 [lines 5 and 6]. The analyst assesses that if the division is in category 1, the minimum percentage of full-strength equipment that it would have is 70 percent, so he enters 0.7 [line 7]. The program then asks him for the maximum value of EQUIP given that category is in state 1 [line 8], and he enters 100 percent [line 9].

At this point, the program displays an initial likelihood distribution for EQUIP given that the division is in category 1 [Figure 3-22]. The analyst believes this distribution is a good representation of the true distribution, at least as a first pass, so he selects RETURN with the light pen. The program next asks the analyst for the minimum value [Figure 3-21, lines 10 and 11] and maximum value [line 13] of EQUIP when CAT is in state 2, and the analyst specifies the values 40 percent [line 12] and 100 percent [line 14]. As before, the program displays a symmetrical curve between these two values [Figure 3-23], but since the analyst does not believe the true distribution is symmetrical, he decides to change it. He believes that, although it is possible that the division would have 100 percent of its full-strength equipment if it is in category 2, it is much more likely to have somewhere in the range of 60 percent of its full-strength equipment. This means that the bulk of the distribution should be over the left part of the range, so the analyst touches the circles on the

NORMALIZE
FILE DISTRIBUTION

		UPPER NODE CAT		
		1	2	3
LOWER DIST	1	.750	.500	.150
	2	.250	.167	.250

Figure 3-18

NORMALIZE
FILE DISTRIBUTION

		UPPER NODE CAT		
		1	2	3
LOWER DIST	1	.750	.750	.150
	2	.250	.250	.850

Figure 3-19




Figure 3-20

- 1 THE UPPER NODE IS "CAT "
- 2 IT HAS 3 STATES
- 3 THE LOWER NODE IS "EQUIP "
- 4 IT IS CONTINUOUS
- 5 WHAT IS THE MINIMUM VALUE OF "EQUIP "
- 6 WHEN "CAT " IS IN STATE 1?
- 7 0.7
- 8 WHAT IS THE MAXIMUM VALUE OF "EQUIP "
- 9 1.0
- 10 WHAT IS THE MINIMUM VALUE OF "EQUIP "
- 11 WHEN "CAT " IS IN STATE 2?
- 12 0.4
- 13 WHAT IS THE MAXIMUM VALUE OF "EQUIP "?
- 14 1.0
- 15 WHAT IS THE MINIMUM VALUE OF "EQUIP "
- 16 WHEN "CAT " IS IN STATE 3?
- 17 0.15
- 18 WHAT IS THE MAXIMUM VALUE OF "EQUIP "?
- 19 1.0

Figure 3-21




Figure 3-22




Figure 3-23

display with his light pen and moves them, one at a time, toward the left [Figure 3-24]. When he has moved all of the circles to their desired location, he touches MOVE CURVE with the light pen to have the program draw the distribution through the circles [Figure 3-25]. Satisfied for the moment, the analyst hits RETURN with the light pen.

The program continues with the assessment by asking the analyst for the extremes of the distribution over EQUIP when the CAT variable is in state 3 [Figure 3-21, lines 15, 16, and 18]. When he has provided them [lines 17 and 19], the program displays the initial distribution [Figure 3-26]. As with the preceding distribution, the analyst believes that the bulk of the distribution ought to lie in the left portion of the range, so he moves the circles to the left [Figure 3-27]. He hits MOVE CURVE, and the program moves the curve [Figure 3-28]. The analyst sees that he has inadvertently put a lump in the curve, so he moves the appropriate circle down slightly [Figure 3-29], selects MOVE CURVE, and the program draws the curve through the circles [Figure 3-30]. The analyst hits RETURN with the light pen, and the program, having processed all of the required distributions, displays all three distributions on the same graph [Figure 3-31].

Upon inspection of the display, the analyst realizes that he has made the distribution for category 3 too uniform; there should be a higher probability that the percentage of full-strength equipment will lie in the range around 40 percent and a lower probability that the percentage will be in the upper range. Accordingly, he touches the light pen to the curve marked 3, and the program responds by displaying the distribution for state 3 [Figure 3-32]. The analyst moves several of the circles toward the mode of the distribution to describe a narrower distribution [Figure 3-33] and then touches MOVE CURVE with the light pen [Figure 3-34]. When he selects RETURN, the program returns to the composite display [Figure 3-35].

The analyst also believes that the distribution for category 2 is a bit too broad, so he selects curve 2 with the light pen, and the program displays the distribution for category 2 [Figure 3-36]. The analyst adjusts the circles [Figure 3-37], selects MOVE CURVE [Figure 3-38], and returns to the composite display [Figure 3-39]. At this point, the analyst is satisfied with the complete likelihood distribution, so he selects FILE DISTRIBUTION, and the program returns to the tree structure display [Figure 3-40].

3.3 CASE 3: Continuous Upper Variable, Discrete Lower Variable

The analyst next decides to assess the distribution relating PERS and CAT. The program begins by verifying the identity of the nodes selected by the analyst and the number of




Figure 3-24




Figure 3-25




Figure 3-26




Figure 3-27




Figure 3-28




Figure 3-29




Figure 3-30




Figure 3-31




Figure 3-32




Figure 3-33




Figure 3-34




Figure 3-35




Figure 3-36




Figure 3-37




Figure 3-38




Figure 3-39




Figure 3-40

states associated with them [Figure 3-41, lines 1 through 4]. Then the program asks the analyst for the minimumn value of the upper-level variable, PERS [line 5], and the analyst types 0.0 [line 6]. In order to define the distribution at this point, the program presents the analyst with a probability histogram with which to make his assessment [Figure 3-42]. Since the lower-level variable, CAT, has three states, there are three histogram bars on the display. To make his assessments, the analyst adjusts the heights of these bars with the light pen to indicate the probability of each of the three states of CAT given that PERS has a value of 0 (that is, there are no personnel in the division). Under these conditions, the analyst is certain that the division is in readiness category 3, so he moves the first two histogram bars down to zero with the light pen [Figure 3-43], and then selects NORMALIZE with the light pen. The program normalizes the probability histogram so that the probabilities sum to 1.0 [Figure 3-44]. The analyst then selects RETURN, and the program moves on to the next assessment.

The program defines the next assessment at the maximum value for PERS. The program requests this value [Figure 3-41, line 7], and the analyst enters 10,000 [line 8]. As before, the program presents the analyst with the probability histogram [Figure 3-45], and the analyst moves the second and third histogram bars down to zero [Figure 3-46], since if the division were known to have 10,000 personnel, it would be virtually certain to be a category 1 division. The analyst normalizes the distribution as before [Figure 3-47] and touches RETURN.

In order to describe the entire likelihood function, the analyst must now supply likelihood distributions for several intermediate values of the upper-level variable. He is free to use any values between the minimum and maximum values that he wishes, and in any order, provided that two assessments are not too close together (the program will warn him if this situation occurs). The program asks him for the value of the upper-level variable he wishes to assume hypothetically to be true [Figure 3-41, line 9], and the analyst replies by typing 3,000 [line 10]. The analyst is given the histogram [Figure 3-48] and moves the bars so that the most likely state of CAT given that PERS is equal to 3,000 is category 3. In fact, the analyst believes that there is about a 10 percent chance that the division might be in readiness category 2 [Figure 3-49]. The analyst selects RETURN, and the program displays a new menu asking if he wishes to make further assessments [Figure 3-50]. The analyst selects YES. [Note: The analyst does not need to normalize the distribution before selecting RETURN: the program will do so automatically.]

1 THE UPPER NODE IS "PERS."
2 IT IS CONTINUOUS.

3 THE LOWER NODE IS "CAT."
4 IT HAS 3 STATES.

5 WHAT IS THE MINIMUM VALUE OF "PERS"?
6 0.0

7 WHAT IS THE MAXIMUM VALUE OF "PERS"?
8 10000.0

9 WHAT IS THE CURRENT VALUE OF "PERS"?
10 3000.0

11 WHAT IS THE CURRENT VALUE OF "PERS"?
12 4500.0

13 WHAT IS THE CURRENT VALUE OF "PERS"?
14 5000.0

15 WHAT IS THE CURRENT VALUE OF "PERS"?
16 5500.0

17 WHAT IS THE CURRENT VALUE OF "PERS"?
18 6000.0

19 WHAT IS THE CURRENT VALUE OF "PERS"?
20 6500.0

21 WHAT IS THE CURRENT VALUE OF "PERS"?
22 7000.0

23 WHAT IS THE CURRENT VALUE OF "PERS"?
24 7500.0

25 WHAT IS THE CURRENT VALUE OF "PERS"?
26 8000.0

27 WHAT IS THE CURRENT VALUE OF "PERS"?
28 4000.0

Figure 3-41




Figure 3-42




Figure 3-43




Figure 3-44




Figure 3-45




Figure 3-46




Figure 3-47




Figure 3-48




Figure 3-49




Figure 3-50




Figure 3-51

The analyst specifies that the next value of PERS is 4,500 [Figure 3-41, lines 11 and 12], the program displays a new histogram [Figure 3-51], and the analyst moves the probability histograms [Figure 3-52]. In this case, he believes that a division with a personnel strength of 4,500 has about a 60 percent chance of being a category 2 division, about a 40 percent chance of being a category 3 division, and almost no chance of being a category 1 division. The analyst selects RETURN, is asked whether he wishes to make more assessments [Figure 3-53], and selects YES.

In this way the analyst specifies the probabilities of categories 1, 2, and 3 for several values of personnel [Figure 3-41, lines 13-26]. When he has specified all of the distributions that he wants to, he merely selects NO when presented with the ANOTHER ASSESSMENT? menu. Then the program displays the resulting complete likelihood distribution [Figure 3-54]. This distribution represents the probability that the division is in category 1, or 2, or 3 given any particular value of the upper-level variable, personnel strength. The vertical bars on the display represent the values of the upper-level variable which the analyst specified in making his assessments. Thus, they represent the probability distributions which actually determine the shape of the curves. Because the analyst may have been in error on his original assessments, or because they may have been incomplete, he is given an opportunity at this point to add a new assessment, remove an old assessment, or change the probabilities associated with one of the old assessments.

After looking at the entire likelihood distribution, the analyst decides that he has assigned too much probability to the possibility of a division with small numbers of personnel being a category 2 division. To remedy this, he decides to move the first assessment he made, at a personnel strength of 3,000, up to 4,000. To make this change, he must first delete the assessment at 3,000. He first touches REMOVE ASSESSMENT, then touches the vertical bar located at 3,000. The program deletes the assessment at 3,000 personnel, and redraws the curves using the remaining assessments [Figure 3-55]. Next the analyst adds the desired assessment at 4,000 personnel. He touches ADD ASSESSMENT and the program types WHAT IS THE CURRENT VALUE OF "PERS"? [Figure 3-41, line 27], and the analyst specifies the value 4,000 [line 28]. He is given the probability histogram [Figure 3-56] and makes his assessment as before [Figure 3-57]. The analyst hits RETURN, and the complete display is drawn, incorporating the new assessment [Figure 3-58].

Upon inspection of this display, the analyst decides to make one minor change. He believes that the probability of




Figure 3-52




Figure 3-53




Figure 3-54




Figure 3-55




Figure 3-56




Figure 3-57




Figure 3-58




Figure 3-59

category 2 around 4,500 to 5,000 personnel should be a little steeper, so he wishes to alter his assessment at the point 4,500. In order to do this, he only has to touch the appropriate assessment bar, and the program will display the probability histogram for that assessment, showing his previously assessed values [Figure 3-59]. Upon consideration of this new assessment, the analyst decides that a division which has 4,500 personnel is equally likely to be either a category 2 or a category 3 division. He touches the light pen to the bar over category 2 and lowers it until it has the same height as the bar over category 3 [Figure 3-60], touches NORMALIZE [Figure 3-61], and RETURN. The program displays the new likelihood distribution [Figure 3-62]. At this point the analyst is satisfied with the likelihood distribution, so he touches FILE DISTRIBUTION with the light pen. The program displays the likelihood distribution without the assessment bars [Figure 3-63], files the distribution, and returns to the display of the tree structure [Figure 3-64].

3.4 CASE 4: Continuous Upper and Lower Variables

The next assessment which the analyst wishes to make is the likelihood of various barracks capacities given possible personnel strengths. However, the likelihood distribution which links PERS with BARCAP is really a likelihood surface, since both variables are continuous. This likelihood surface is approximated in the procedure by a number of distributions oriented in one direction; that is, a set of likelihood distributions showing the likelihood of each of the possible values of the lower-level variable given particular values of the upper-level variable. These distributions are used by the program to fill in the complete distribution surface.

The analyst touches ASSESS, then the connection between BARCAP and PERS. The program first verifies the identity of the connected variables [Figure 3-65, lines 1 through 4], then requests the minimum and maximum values of the upper-level variable [lines 5 and 7]. The analyst believes personnel strength for this division to be somewhere in the range from 0 to 10,000, so he enters these values [lines 6 and 8]. The program next asks the analyst for the current value of PERS, that is, the value currently assumed to be true hypothetically [line 9]. The analyst may give these values in any order, and they need not be complete, since he will be given an opportunity to change, add, or delete assessments later. The analyst begins by hypothetically assuming that personnel strength is known to be 1,000 [line 10]. The program then asks the analyst for the conditional minimum and maximum values of BARCAP [lines 11, 12, and 14]. The reason that the program requests minimum and maximum values for each conditioning value of the upper-level variable is that, for any given value of the upper-level variable, a large portion of the total range of values may have




Figure 3-60




Figure 3-61




Figure 3-62




Figure 3-63




Figure 3-64

- 1 THE UPPER NODE IS "PERS".
- 2 IT IS CONTINUOUS.

- 3 THE LOWER NODE IS "BARCAP".
- 4 IT IS CONTINUOUS.

- 5 WHAT IS THE MINIMUM VALUE OF "PERS"?
- 6 0.0

- 7 WHAT IS THE MAXIMUM VALUE OF "PERS"?
- 8 10000.0

- 9 WHAT IS THE CURRENT VALUE OF "PERS"?
- 10 1000.0

- 11 WHAT IS THE MINIMUM VALUE OF "BARCAP"?
- 12 WHEN "PERS" IS EQUAL TO 1000.00000?
- 13 1000.0

- 14 WHAT IS THE MAXIMUM VALUE OF "BARCAP"?
- 15 8400.0

- 16 WHAT IS THE CURRENT VALUE OF "PERS"?
- 17 2000.0

- 18 WHAT IS THE MINIMUM VALUE OF "BARCAP"?
- 19 WHEN "PERS" IS EQUAL TO 2000.00000?
- 20 2000.0

- 21 WHAT IS THE MAXIMUM VALUE OF "BARCAP"?
- 22 8700.0

- 23 WHAT IS THE CURRENT VALUE OF "PERS"?
- 24 3000.0

- 25 WHAT IS THE MINIMUM VALUE OF "BARCAP"?
- 26 WHEN "PERS" IS EQUAL TO 3000.00000?
- 27 3000.0

- 28 WHAT IS THE MAXIMUM VALUE OF "BARCAP"?
- 29 9100.0

- 30 WHAT IS THE CURRENT VALUE OF "PERS"?
- 31 4000.0

- 32 WHAT IS THE MINIMUM VALUE OF "BARCAP"?
- 33 WHEN "PERS" IS EQUAL TO 4000.00000?
- 34 4000.0

- 35 WHAT IS THE MAXIMUM VALUE OF "BARCAP"?
- 36 9500.0

Figure 3-65

37 WHAT IS THE CURRENT VALUE OF "PERS" ?
38 5000.0

39 WHAT IS THE MINIMUM VALUE OF "BARCAF" ?
40 WHEN "PERS" IS EQUAL TO 5000.00000?
41 5000.0

42 WHAT IS THE MAXIMUM VALUE OF "BARCAF" ?
43 9800.0

44 WHAT IS THE CURRENT VALUE OF "PERS" ?
45 6000.0

46 WHAT IS THE MINIMUM VALUE OF "BARCAF" ?
47 WHEN "PERS" IS EQUAL TO 6000.00000?
48 6000.0

49 WHAT IS THE MAXIMUM VALUE OF "BARCAF" ?
50 10200.0

51 WHAT IS THE CURRENT VALUE OF "PERS" ?
52 7000.0

53 WHAT IS THE MINIMUM VALUE OF "BARCAF" ?
54 WHEN "PERS" IS EQUAL TO 7000.00000?
55 7000.0

56 WHAT IS THE MAXIMUM VALUE OF "BARCAF" ?
57 10500.0

58 WHAT IS THE CURRENT VALUE OF "PERS" ?
59 8000.0

60 WHAT IS THE MINIMUM VALUE OF "BARCAF" ?
61 WHEN "PERS" IS EQUAL TO 8000.00000?
62 8000.0

63 WHAT IS THE MAXIMUM VALUE OF "BARCAF" ?
64 10800.0

65 WHAT IS THE CURRENT VALUE OF "PERS" ?
66 9000.0

67 WHAT IS THE MINIMUM VALUE OF "BARCAF" ?
68 WHEN "PERS" IS EQUAL TO 9000.00000?
69 9000.0

70 WHAT IS THE MAXIMUM VALUE OF "BARCAF" ?
71 11000.0

Figure 3-65 (Continued)

a very small probability. In such a case, it would be difficult for the user to describe the shape of the curve in the small range having a significant probability.

The analyst specifies that if personnel strength is known to be 1,000 men, then barracks capacity must be greater than 1,000 and could be as high as 8,400 [lines 13 and 15]. At this point, the program draws a symmetrical distribution between the minimum and maximum values of BARCAP [Figure 3-66]. (NOTE: This is exactly the same display that the analyst encountered in CASE 2.) The analyst begins specifying the shape of the distribution by moving the circles. In this case, the analyst believes that it is much more likely that the barracks capacity will be nearer to 1,000 men than to 8,400 men if the true personnel strength is equal to 1,000 men, so he moves the circles to the left [Figure 3-67]. He then touches MOVE CURVE with the light pen, and the program shifts the curve to fall through the circles [Figure 3-68]. He is satisfied with the distribution, so he touches RETURN with the light pen, and is presented with a new menu [Figure 3-69]. He wishes to make more assessments, so he touches YES.

The program requests information as before. It first asks for the value of the upper-level variable currently assumed to be true [Figure 3-65, line 16]. The analyst specifies 2,000 [line 17], the program requests the maximum and minimum values of BARCAP [lines 18, 19, and 21], and the analyst specifies the range 2,000 to 8,700 [lines 20 and 22]. The program displays the curve [Figure 3-70], the analyst moves the circles to the left [Figure 3-71] and touches MOVE CURVE. The program moves the curve [Figure 3-72], the analyst touches RETURN and is asked if he wishes to make more assessments [Figure 3-73]. He touches YES.

The analyst proceeds in this fashion to specify a number of distributions [Figure 3-65, lines 23 through 71 and Figures 3-74, 3-75, 3-76, 3-77, 3-78, 3-79, and 3-80]. Note that these distributions differ only slightly in shape, but the ranges associated with the distributions are quite different. Thus, the normalized shapes of the distributions will be quite different. When he has reached the last distribution he wishes to assess and is asked whether he wishes to make another [Figure 3-81], the analyst simply touches NO with the light pen. Having completed, at least initially, the set of assessments for this distribution, the program displays the complete likelihood distribution [Figure 3-82]. Here the x-axis represents the upper-level variable, PERS, and the z-axis (extending into the background) represents the lower-level variable, BARCAP. Thus, each curve represents the probability over the possible values of BARCAP given the particular value of PERS determined by the point at which the curve touches the x-axis.




Figure 3-66




Figure 3-67




Figure 3-68




Figure 3-69




Figure 3-70




Figure 3-71




Figure 3-72




Figure 3-73




Figure 3-74




Figure 3-75




Figure 3-76




Figure 3-77




Figure 3-78




Figure 3-79




Figure 3-80




Figure 3-81




Figure 3-82




Figure 3-83

At this point, the analyst may add more assessments, delete assessments, or change assessments he has already made in exactly the same fashion as he did in CASE 2. To change an assessment, he only needs to touch the curve he wishes to change with the light pen. To add an assessment, he touches ADD ASSESSMENT with the light pen, and the program requests the necessary information about the current value of PERS and the minimum and maximum values of BARCAP, and displays the assessment curve. To delete an assessment, the analyst touches REMOVE ASSESSMENT, and then touches the curve for the assessment he wishes to delete. In this case, the analyst is satisfied with all of his assessments, so he touches FILE DISTRIBUTION and returns to the structure display [Figure 3-83].

4.0 PROCESSING THE MODEL

Now the analyst wishes to process the model to obtain the composite likelihood distribution. Before processing, however, he wishes to check whether he had previously assessed the likelihood distribution relating barracks capacity and barracks area. He touches EXAMINE, and then touches the arrow connecting AREA to BARCAP. The program responds by finding the distribution (if it exists) and displaying it. In this case the distribution does exist so the program displays it [Figure 4-1].

Having satisfied himself that the correct distribution is present, the analyst touches RETURN, and the program returns to the tree structure [Figure 4-2]. He then touches PROCESS with the light pen, and the program begins processing starting with the data nodes. In this example, the only data node which has more than one state is the DIST node. The program types DATA NODE "DIST" HAS MORE THAN ONE STATE. WHICH STATE HAS BEEN OBSERVED? [Figure 4-3, lines 1-2]. The analyst knows that the division is near the border, so he types "1" for the state of the variable DIST [line 3].

The program processes the remainder of the tree, finding the distributions it requires, and produces the composite likelihood distribution [Figure 4-4]. In this example, it has been assumed that the analyst had previously assessed most of the distributions prior to the current session. If any of the distributions were missing or were assessed in a manner inconsistent with the present model (such as having the wrong number of states associated with a variable), the program would have halted at the point of the error and informed the analyst of the location in the tree of the missing or inconsistent data. Then the analyst could either assess the required distribution or save the model for use at a time when he wished to supply the missing distribution.

The composite likelihood distribution shows the relative likelihood that any particular state of the upper-level variable (in this case, PERS) would have produced the observed data. In order to infer the probable strength of the division, the analyst would combine this composite likelihood distribution with the prior distribution (the probability distribution representing the analyst's belief about the strength of the division prior to receiving any specific information) by using Bayes' Theorem. This would give him the posterior distribution (the distribution over personnel strength based upon all the data).




Figure 4-1




Figure 4-2

- 1 DATA NODE "DIST." HAS MORE THAN ONE STATE
- 2 WHICH STATE HAS BEEN OBSERVED?
- 3 1

Figure 4-3




Figure 4-4

The composite distribution shows that the most likely personnel strength to have produced the observed data is around 7,500 men, and any division with a personnel strength of less than 5,000 men or greater than 9,000 men is extremely unlikely to have produced the data. In addition to the composite likelihood distribution, the analyst may also inspect specific portions of the model in order to discover how different factors in the model contributed to the final composite likelihood distribution using the EXAMINE option. For example, the analyst is interested in the effect that the observed barracks area, OBS, had on the final likelihood distribution for personnel strength, PERS. He first touches RETURN, the program displays the tree structure [Figure 4-5], then the analyst touches EXAMINE, and the arrow connects PERS to BARCAP. The program displays the relative likelihood for all the data below PERS (in this case, the only data node is OBS), given each possible value of PERS [Figure 4-6]. He can see that the information about observed barracks area gives only a very vague idea of the personnel strength, since the data must pass through two intervening variables: true barracks area (AREA) and barracks capacity (BARCAP). The uncertainty associated with each of these nodes results in a diffuse distribution at PERS. The analyst wishes to return to the main diagram, so he touches RETURN (Figure 4-7).

The analyst is also curious about the contribution made to the likelihood distribution over category by the observed equipment, COUNT. He touches EXAMINE, then touches the arrow connecting EQUIP to CAT. The program displays the likelihood distribution of all data below EQUIP (in this case just COUNT) given each possible state of CAT [Figure 4-8]. Since CAT is a discrete variable, the likelihood distribution is simply a row vector, showing the relative likelihood that each type of category would produce the observed equipment count. Thus, the most likely readiness category, if one depended solely on equipment count, would be category 1. The analyst touches RETURN, and the program returns to the tree display [Figure 4-9].

This completes all the assessments which the analyst wishes to make by using the model at this time. However, because he wishes to make further use of the model in the future, he touches SAVE MODEL with the light pen, and the program saves the model. To conclude the session, he touches STOP.




Figure 4-5




Figure 4-6




Figure 4-7




Figure 4-8




Figure 4-9